We thank all the reviewers for their detailed and insightful feedback.

Reviewer 1:

We would like to clarify that our setting does include "reservation price" pricing policies; in fact, they are the principal focus of our paper. Specifically, the goal of our paper is to learn a good linear reservation pricing policy; a policy which examines the features of the item for sale, generates a reservation price from those features, and runs a second-price auction with this price as a reserve (in the case of a single bidder, this reduces to posting a single price for the item). See Figure 1 in the appendix. The 1-Lipshitzness constraint is not a constraint on the function determining how much a bidder pays given their bid (which indeed would eliminate many natural mechanisms), but rather a constraint on the function determining the reservation price as a function of the features.

In particular, note that our benchmark contains the optimal mechanism in the simple (non-contextual) setting.

Reviewer 2:

We thank the review for the suggestions to improve the paper which we plan to implement in the revision: we will extend the discussion on motivation and add more context to the results. Below we respond to questions asked:

- Why is hardness surprising? Indeed the discontinuity of the loss function hints at hardness, but it has been open for a while and in fact we are not aware of any straightforward reduction. Even though hardness is not surprising it wasn't at all clear where the problem was situated in the hardness-spectrum,i.e. it is APX-hard or constant approx or PTAS or FPTAS?
- Constants in the proofs: constants are somewhat arbitrary as we don't try to optimize them (instead we simply pick constants that are convenient for the analysis). We choose to write explicit constants that satisfy the necessary inequalities (instead of just using $O(\cdot)$) since the conditions for both JL and the hardness reduction are somewhat delicate and often require checking both upper and lower bounds.
- We discuss the loss of [MM14] in the last paragraph of Section 4 in our paper. We call it the SPA loss. It takes two bids (v_1, v_2) and returns the revenue of a second price auction with reserve. Throughout our paper we study the simpler pricing loss which is equivalent to SPA with $v_2 = 0$. However in the end of Section 4 we note that all our results extend to SPA and any other lower-semi-Lipschitz loss function.
- The question of consistent surrogates is an interesting one. The existence of convex consistent surrogates is ruled-out by our hardness result. One could still search for consistent surrogates without poly-time guarantees but that are smooth enough to be tractable in practice. Since our focus here was to obtain provable guarantees, we opted for working with the loss function directly.

Reviewer 3:

Thanks for the suggestions to improve the paper. We plan to implement those in the revision. We will add a discussion in future directions on how to extend this work to non-truthful auctions (such as first-price and GSP) as well as more motivation for the truthful case (e.g single slot second price and VCG auctions often used in display ads). We will expand the discussion on the future direction of strategic response in the conclusion.